Reference Price Shifts and Customer Antagonism: Evidence from Reviews for Online Auctions

Tobias Gesche*

Harvard University & University College London

*SECONDARY JOB MARKET PAPER

October 31, 2017

Abstract

Based on data from a large-scale sales campaign on eBay, I show that successful auction customers punish the seller through unfavorable public feedback when they later learn about a cheaper fixed-price offer. The probability of such feedback is four times as big for auction customers as for customers who bought the same item from the same seller for the fixed price. Remarkably, this probability is increasing in the auction price, even though auction customers actively shaped this price themselves. This price effect on unfavorable feedback is concentrated in a period immediately after the auction. In this period, customers could learn about the fixed-price offer but could not gather any new information regarding the item they had just acquired because it had not yet been shipped. These findings are in line with a model of feedback-giving which combines reciprocal with reference-dependent preferences.

Keywords: customer antagonism, pricing, reference prices, online reputation, eBay

JEL Classification: D44, D91, M31

Contact: tobias.gesche.job@gmail.com

I would like to thank to Réné Algesheimer, Björn Bartling, Alessandro Bonatti, Gregory Crawford, Brent Hickman, Michel Maréchal, Frederick Schwerter, and Duncan Simester for helpful comments and discussions. I also thank Andreas Roider for encouragement during the very early stages of this project and acknowledge support by the Swiss National Science Foundation (SNSF Grant #P2ZHP1_171900).

1 Introduction

Pricing is crucial for sellers and policymakers alike for reasons that go beyond the resulting allocations and transfers. Longstanding evidence shows that not only the price itself but also the circumstances under which it is determined affect how customers evaluate a transaction (Kahneman et al., 1986; Frey and Pommerehne, 1993; Xia et al., 2004). Negative feelings about a transaction can even lead to concrete actions by customers against sellers, for example when some are charged different prices than others (Yi and Baumgartner, 2004; Anderson and Simester, 2008, 2010). This is particularly relevant for online markets as their flexible nature allows to vary sales conditions and prices across customers, either as means of experimentation (Einav et al., 2015) or user-based price discrimination (Mikians et al., 2012; Shiller, 2014).[1] Auctions could, in principle, be a solution. They are easily implemented online and allow differential pricing. At the same time, they have the potential to prevent dissatisfied customers who are not passive price-takers but, through their bids, are consciously and willingly determining the prices they pay (Chandran and Morwitz, 2005; Hinz et al., 2011). However, this paper's findings show that when auctions co-exist with fixed-price offers, they cause strong customer antagonism through typical features of online markets – reputation systems and customers' limited perception of competing offers.

Using data from a large-scale online sales campaign, I report on the determinants of post-sale behavior of auction customers towards the seller. During the campaign, the seller first used an auction to sell several thousand units of an item. Two days later, the same seller (a railway company) sold the same item (a voucher for an open-destination rail journey) on the same sales platform (eBay) for a fixed price. I find that customers who bought the item via the auction are four times as likely to use the market platform's reputation system to give the seller an unfavorable feedback as customers who bought the item for a fixed price. I also find that the more the auction customers paid, the more likely they are to punish the seller through unfavorable feedback. This finding is hard to reconcile with the fact that the customers themselves played a crucial part in determining the auction price. They could have easily prevented to pay a price over which they later become antagonized by bidding accordingly. To explain their behavior, I develop a simple model of feedback-giving in which reference prices matter. The key effect identified in this model is that a downward-shift in reference prices – as caused by observing a

[1] A well-known example where this backfired is Amazon's attempt to charge regular customers higher prices than new ones which lead to pronounced criticism when discovered (see Ward, 2000). Turow et al. (2005) provides survey evidence that American internet shoppers are largely unaware of how personal information is used in online retail but that they condemn its use for (differential) pricing when presented with such scenarios. For a review of data-driven differential pricing and its legal challenges, see the recent report by the White House's Council of Economic Advisers (CEA, 2015).

lower fixed-price offer after the auction ended – can lead customers to negatively review a transaction which, had such a shift not occurred, would not have yielded an unfavorable review.

In line with this theory, I show that the negative effect of a higher auction price on feedback is exclusively concentrated in the period immediately after the auction ended but before the obtained item arrived by post. In this period, successful bidders could get no new information regarding a specific transaction. However, through various channels, they could learn about the fixed-price sale which occurred shortly after the auction ended. Accordingly, the only new information which customers got during this initial period and which could have affected feedback was with respect to reference prices. Right afterwards, when items started to arrive and customers' actual user experience could start to determine feedback, the price effect on adverse feedback drops to zero and stays flat.

The setting considered here allowed auction customers to know about the cheaper fixed-price offer before they bid. However, their feedback's timing and the price effect indicate that many customers just followed the advertisement for the auction and overlooked other upcoming offers. When they then discover, ex-post relative to their transaction, that they missed a better deal, they retaliate. These results therefore demonstrate how the combination of customers' limited attention and differential pricing can backfire on the seller's reputation. They also show that this is caused by reference price shifts, i.e. features which are unrelated to the idiosyncratic, objective characteristics of single transactions.

Related literature: This paper connects to and extends several branches of the literature. First and foremost, it relates to the prior research on fair pricing and consumer antagonism. In a classic study, Kahneman et al. (1986) surveyed Canadian respondents on what they consider acceptable pricing practices in several hypothetical situations. Their findings indicate that a general notion of non-exploitation is important. In particular, exploiting customers via differential pricing schemes is considered unacceptable if it serves to increase seller profits. Similar findings, in particular with respect to customers' aversion to differential pricing, were also obtained in a subsequent study with German and Swiss respondents (Frey and Pommerehne, 1993) and other surveys (for a review, see Xia et al., 2004). Using data from online purchases, this paper confirms these finding and adds new insights. In particular, I show that even when customers selected themselves into a price-differentiating sales mechanism such as an *actual* auction (i.e. not a hypothetical situation during a survey) and co-determine its final price, they end up punishing the seller for the outcome. I also show that such punishment is caused by the interaction of customers learning about an alternative offer and the price premium they paid relative to this alternative.

The unfavorable feedback I document is a form of customer antagonism, i.e. customers who convert negative emotions into concrete actions against the seller (Yi and Baumgartner, 2004). Several theoretical accounts have explored the constraints which customer antagonism imposes on sellers' differential pricing strategies (Rotemberg, 2011; DiTella and Dubra, 2014; Battigalli et al., 2015). Recent experimental evidence which is line with such constraints comes from Leibbrandt (2016): Using a laboratory market for real items, he shows that customers cease to buy from sellers who are price-discriminating, even when such a transaction would leave them a surplus. He also finds that sellers anticipate this and limit price-discrimination when they expect buyers to be aware of it.

Anderson and Simester (2008) present field evidence for customer antagonism. They look on a mail-order who charged higher prices for larger cloth sizes. Customers reacted adversely to such a pricing policy by stopping to order from the company, at a magnitude twice as large as the pure price effect would imply. Another field experiment by Anderson and Simester (2010) shows that *lowering* prices can also lead to less demand by antagonized customers. In this study, they randomized price discounts in a new edition of another mail-order's catalog. Their results show that customers who had bought the reduced item before, i.e. at a higher price, subsequently ceased to order from the mail order. This effect is most concentrated among the customers who had previously paid the most, relative to the discount.

Besides providing a model which accommodates these results, this paper adds to the empirical literature on customer antagonism along three main dimensions: First, I demonstrate the relevance of customer antagonism in *online* retail, a large and steadily growing market. Second, I show that it can not only manifest through customers who cease to buy from a seller but also through attacks on the seller's reputation. This is particularly relevant for online markets which crucially rely on accurate feedback and reputation systems (Dellarocas, 2003; Tadelis, 2016). Third, this is, to my knowledge, the first study which demonstrates that antagonism can also arise in auctions, i.e. when customers themselves are not passive price takers but play a crucial part in the price-setting process.

The results presented here also relate to the research on how experience shapes reference points and the economic consequences of this, for example in the context of contract re-negotiation (Hart and Moore, 2008; Fehr et al., 2011), relative price perceptions (Simonsohn and Loewenstein, 2006; Weaver and Frederick, 2012; Bordalo et al., 2017) or the effect of auctions' start and reserve prices on bidding behavior (Ariely and Simonson, 2003; Kamins et al., 2004). Closely related is also Herz and Taubinsky (2017). In their experiment, subjects first experienced ultimatum games with either responder or proposer competition which lead to relatively low or high offers, respectively. These subjects' active and passive experiences in the initial games affected what they later consider as fair bargaining outcomes, i.e. the amounts they offer or accept in subsequent ultimatum games without competition. I also show that reference prices are updated, based on lower prices offered to others and that this can trigger negative emotions. In this regard, my findings are similar to Card and Dahl (2011) who link increased domestic violence to unexpected losses in football games and Mas (2003) who finds higher crime reports and lower arrests after police unions lost wage arbitrations. However, the hostile behavior documented here occurs in a different setting, a highly organized virtual market place. It also indicates another mechanism in how reference points can cause negative feelings and actions: Instead of being caused by a sudden downward shift in outcomes for a given reference point, the negative actions which I document are caused by a downward shift in reference prices for a given transaction outcome.

This paper's findings also show that successful bidders were often surprised to learn, after the auction had ended, about the cheaper fixed-price offer – even though they could have figured out about these offers before they bid. Therefore, this work also links to the literature on limited attention and under-searching in online environments where offers are easier searchable but also more vast and differentiated than offline (Brynjolfsson et al., 2011). For example, online customers frequently rely on salient cues such as prominent digits of used cars' odometers and differences in these cars' first registration years rather than absolute age differences (Lacetera et al., 2012; Englmaier et al., 2016). They also often neglect extra fees (Hossain and Morgan, 2006) and herd with other bidders (Simonsohn and Ariely, 2008). For online auctions, Ariely and Simonson (2003) and Malmendier and Lee (2011) show that customers often bid more than what is necessary to obtain the same item via a fixed-price sale. While there is some discussion whether this is due to limited attention or too high search costs and whether this ought to be called over-bidding (Schneider, 2016; Malmendier, 2016), the fact that alternative, cheaper offers are left unused is undisputed. This paper's findings show that this does not only harm a customer who misses a better offer but also the seller if the customer later finds out.

The next section develops a simple model of feedback-giving and how reference price shifts influence it. Section 3 presents the dataset analyzed here and the setting from which it was obtained. Based on this description, Section 4 relates the data to the model by deriving predictions and an identifying strategy; the corresponding results are presented in Section 5. Section 6 concludes and discusses the findings' managerial and economic implications with regards to the role of online reputation systems, sellers' pricing strategies, and the role of auctions in retail. An appendix collects further material.

2 Model

In the following, I will set up a simple model which combines reciprocal with reference-dependent preferences. The model describes how a customer ("she") evaluates her purchases and how this influences her behavior towards the seller ("he"). It takes an ex-post perspective by looking at how, given a customer's purchase decision, subsequent changes in her reference point affect her actions. Based on this analysis, I will then derive predictions which will guide the empirical analysis.

Consider a customer who has obtained an item at period $t = 0$ for a price p. At the time of the purchase, the item has expected value v for the customer. In a later period $t \in \{1, 2, ...\}$, the customer may then get new information about the object which she did not have initially. This information is denoted by the term $e_t \in \mathbb{R}$ and represents (positive or negative) user experience, for example regarding moral hazard in the seller's post-transaction behavior or the item's quality. A customer's valueation at the transaction date reflects this expectation, i.e. $e_0 = 0$ can always be assumed. At some period $\tau > 0$, the user experience realizes. The net utility of the transaction as experienced by the customer in period t is then given by $u_t = v - p + \epsilon_t$ with $\epsilon_t = 0$ if $t < \tau$ and $\epsilon_t = e_\tau$ for $t \geq \tau$.

I also allow a customer to have reference-dependent utility. How she assesses the transaction is therefore not only dependent on her net utility u_t but also how this compares to the reference utility $u_t^r = v - r_t$ of buying the item elsewhere at price r_t. This additional reference-dependent utility is then given by $\mu(u_t - u_t^r) = \mu(r_t - p + \epsilon_t)$ where $\mu \geq 0$ scales this utility in relation to the base net utility u_t. Changes in the reference-dependent utility therefore occur either through user experience ($e_t \neq 0$) and/or through an update in the reference price ($r_t \neq p$). For simplicity, I assume that the initial reference price is the transaction price, thus that $r_0 = p$. I allow for asymmetric reference-dependence by scaling the reference-dependent utility of losses relative to gains with $\lambda > 0$. Loss aversion is then captured by assuming $\lambda > 1$, a parameter range which is also possible here. This would amplify the main effects which will be derived in the following but it is not a necessary assumption. Assuming additivity, a customer's assessment of the transaction at time t is then given by the following expression:[2]

[2]This linear form of reference-dependent utility has been used Köszegi and Rabin (2006) to illustrate applications of reference-dependent utility and in related works that followed (e.g. Heidhues and Köszegi, 2008, 2014). In particular, Lange and Ratan (2010) and Ahmad (2015) use a linear approach model to study how, given reference-dependent preferences, optimal bids in auctions are determined ex-ante. Note that (1) can also be understood as a special case of a more general compound function $\mathcal{A}_t(\{e_k, r_k, v, p\}_{k=0}^t) \equiv \sum_{k=0}^t \alpha_k A_k(e_k, r_k, v, p)$ with time-period specific weights α_k which also takes into account past assessments and which is evaluated at period t. As I will only be interested in the effects which involve contemporary changes, i.e. $\partial \mathcal{A}_t(\{e_k, r_k, v, p\}_{k=0}^t)/\partial x_t = \alpha_t \cdot \partial A_t(e_t, r_t, v, p)/\partial x_t$ with $x_t \in \{e_t, r_t\}$, it is sufficient to focus on the current period t and normalize its weight to one.

$$A_t(e_t, r_t, v, p) = v - p + \epsilon_t + \mu \cdot \big(\max\{r_t - p + \epsilon_t, 0\} + \lambda \cdot \min\{r_t - p + \epsilon_t, 0\}\big) \tag{1}$$

In the following, I will allow customers to take an action which is either in favor of or against the seller, based on this assessment. In the context of this paper, these actions are giving favorable or unfavorable (online) feedback which is why I will henceforth speak of feedback when referring to this action. However, the results apply to any other action with similar consequences. Feedback can be given once for each transaction and is denoted by $f_n \in F$ where F is a discrete, finite and ordered subset of $\mathbb{R}$. I also allow for the possibility that no feedback is given, i.e. no action for or against the seller is taken by the customer. As a convention, I assign an index and value of zero to this case, i.e. $f_0 = 0$ denotes no feedback. Negative elements of F with $n \in \mathbb{Z}_-$ represent an unfavorable (worse than none) feedback while elements of F with $n \in \mathbb{Z}_+$ represent favorable (better than none) feedback. Accordingly, higher positive (negative) values of n denote more favorable (less unfavorable) feedback. I assume that there is always at least one kind of favorable and unfavorable feedback, besides the possibility of giving no feedback, i.e. that $\{f_{-1}, f_0, f_1\} \subseteq F$ holds.

Giving feedback has both, gains and costs to customers. In the context of online feedback, costs of giving feedback can derive, for example, from the time and effort of having to log in to the respective site, search the respective option and writing a comment. These costs of giving feedback f_n are captured by $c(f_n)$ which is the image of a twice continuously differentiable function $c : \mathbb{R} \to \mathbb{R}_0^+$, evaluated at $f_n \in F \subset \mathbb{R}$.[3] Giving no feedback does not create any costs, i.e. $c(0) = 0$. I also assume that c is strictly convex. Therefore, all "actual" feedback ($f_n \neq 0$) is costly and giving more extreme feedback is more costly, for example, because more elaborate wording has to be used for such feedback or because a customer inherently rations stronger statements. Note that c does not need to be symmetric around its minimum. Thus, the costs of giving favorable and unfavorable feedback can grow at different rates, consistent with the findings by Dellarocas and Wood (2008) and Nosko and Tadelis (2015).

Giving feedback also creates utility for a customer, for example through a reciprocity-motive in which punishing (rewarding) a seller for a negatively (positively) assessed transaction yields additional utility. I capture such additional utility of feedback by the term $\psi \cdot (f_n \cdot A_t)$. The variable $\psi > 0$ therefore represents a customer's preference for giving the seller feedback which reflects her assessment relative

[3] This means that only the values of c over F will be relevant. However, defining these costs via a continuous function over the real space simplifies the subsequent exposition.

to the costs of giving feedback.[4] Accordingly, a customer's utility of providing feedback f_n, given her current assessment A_t, is given by

$$U_t\big(f_n|\psi, A_t\big) = A_t(e_t, r_t, v, p) \cdot \big(1 + \psi f_n\big) - c\big(f_n\big). \tag{2}$$

A customer then chooses her feedback so that it maximizes the above expression, i.e. $f_t^* \equiv \arg\max_{f_n \in F} U_t\big(f_n|\psi, A_t\big)$ holds. I will assume that customers are myopic regarding when to issue a non-zero feedback or, equivalently, that they take current perceptions as indicative of future realizations.[5] Therefore, once $f_t^* \neq 0$ holds, they issue the feedback which reflects their contemporary assessment of the transaction. Before and after, they do not issue feedback. Assuming that the motivation to give feedback, as measured by ψ, is heterogeneously distributed across customers according to the strictly increasing c.d.f. $\Psi(x) \equiv \Pr[\psi \leq x]$, one then gets the following:

Proposition 1. *Given an assessment $A_t = A_t(e_t, r_t, v, p)$, it holds that the probability $\Pr[f_t^* \leq f_n|A_t]$ of observing feedback less or equal than some non-maximal feedback score $f_n < \max\{F\}$*

a) is positive and strictly decreasing in A_t for $A_t \neq 0$ and $A_t \cdot f_n > 0$,

b) equals one and is invariant in A_t for $A_t \leq 0 \leq f_n$,

c) equals zero and is invariant in A_t for $A_t \geq 0 > f_n$.

Proof: see Appendix A.

The above result links the distribution of feedback to underlying assessments. Case c) shows that no unfavorable feedback will be issued when the customer's assessment is non-negative. Conversely, case b) implies that no favorable feedback will be issued when the customer's assessment is non-positive.[6] Case a) covers feedback which has the same sign as the underlying assessment and shows that more favorable (unfavorable) feedback is more likely for higher, positive (lower, negative) assessments. Feedback therefore varies with the underlying assessment but only if both are equally signed. In consequence, the comparative statics of $A_t = A_t(e_t, r_t, v, p)$ carry over to equally signed feedback. For unfavorable feedback, this means the following:

[4]Besides reciprocal motives, this formulation also captures complementary altruistic utility of contributing informative feedback to the public good which an unconditional feedback systems effectively is (Avery et al., 1999; Bolton et al., 2004)

[5]This would correspond to $\mathrm{E}[x_\tau] = x_t$ for each $\tau > t$ and $x_t \in \{e_t, r_t\}$ and is consistent with findings that current reference points reflect expectations (see Ericson and Fuster, 2011; Gill and Prowse, 2012; Bartling et al., 2015).

[6]To see this note that case b) implies $\Pr[f_t^* \leq 0|A_t \leq 0] = 1$ and, therefore, $\Pr[f_t^* > 0|A_t \leq 0] = 0$.

Corollary 1. *The probability of observing unfavorable feedback is*

i) price-insensitive if the assessment is positive $\Big(\partial \Pr[f_t^* < 0 | A_t(e_t, r_t, v, p) > 0]/\partial p = 0\Big)$,

ii) increasing in the price if the assessment is negative $\Big(\partial \Pr[f_t^* < 0 | A_t(e_t, r_t, v, p) < 0]/\partial p > 0\Big)$.

Case i) covers situations when there is a positive assessment. A customer will then not issue unfavorable feedback. Accordingly, the price effect with respect to this event is zero. Note that this does not mean that feedback is unaffected by prices. As long as the underlying assessment is positive, a higher price may lead to less pronounced positive positive or even omitted feedback – it is however never negative. Case ii) is relevant when the customer's assessment is negative. In this situation, a higher price paid leads to a lower, negative overall assessment and thereby increases the chance that unfavorable feedback of some given magnitude is issued.

Now assume that the customer has an outside option with a normalized value of zero. Therefore, a transaction is only concluded if at the time of the transaction, the customer's assessment exceeds her outside-option, i.e. when at $A_0 \geq 0$ holds. Given the above assumptions on e_0 and r_0, this means that in a posted offer market, a customer only buys an item with a price that does not exceed her valuation for it. Similarly, in a first- or second-price auctions, a customer's bid is always an upper ceiling on the realized price. A non-negative assessment $A_0 \geq 0$ can then be ensured as long as a bid does not exceed v.[7] Negative feedback and a price-effect as described in part ii) of the corollary then requires a negative assessment and, therefore, a change in the buyer's assessment *after* the transaction is concluded. This can be either due to sufficiently negative experience $e_t < 0$ or due to a sufficiently strong downward revision of a customer's reference price such that $\mu\lambda(r_t - p) < 0$ holds. Note that triggering unfavorable feedback via experience is re-enforced by, but not dependent on, reference-dependent preferences. In contrast, triggering such feedback via references prices is only possible if the customer has sufficiently strong reference-dependent preferences.

The results by Anderson and Simester (2010), which cover posted-offer transactions, are in line with the above model: They document that mail-order customers who received a catalog with discounted items subsequently ceased to order from the mail order company if they had previously purchased the items for a higher price. The action with which customers negatively reciprocate towards the seller is thus not unfavorable feedback but the, in the presence of discounts costly, decision to not order anymore

[7]See the literature review for the role of buyer perception on over-bidding. Over-bidding even with respect to perceived alternatives can also be accommodated, as long as it is not too big, at the costs of a much more complicated exposition.

with the seller. Consistent with an effect induced by reference price shifts, they find that the negative effect of the discount was particularly pronounced for those who had previously ordered at the highest prices relative to the discount, i.e. for whom the $r_t - p < 0$ had the largest magnitude. In the following, I will show a similar effect of how customers react adversely to downward-shifts in their reference prices. In contrast to the findings just described, customer do so via unfavorable feedback and in a context where the original price was not posted by the seller but co-determined by themselves, as the result of an auction.

3 Context and data description

In early August 2008, a large German railway company, in cooperation with the German branch of the internet auction and sales platform eBay, conducted a sales campaign for rail tickets. Starting from August 1 and going till August 10, every day a special offer was available for purchase on a dedicated eBay-page. Each of these daily special offers were vouchers for a train journey which differed from day to day along several dimension, e.g. whether the destination was fixed or free to choose, whether it was for international or national travel or whether the tickets were for first or second-class train rides. Of particular relevance for this paper are the offers of August 1 and 3. On these two days, the offered voucher was the same. It was for a return trip, second class on all domestic trains (except night trains) operated by the railway company. After its sale and payment via money transfer, the paper voucher was shipped by post. In order to use it, customers had to fill in their departure and arrival stations; its specific use was therefore up to the client. Naturally, this results in different valuations for the voucher, depending on the intended use and the opportunity cost of obtaining the ticket elsewhere. At the time of the sale, the regular price for the most expensive itineraries covered by the voucher was 230.00€.

While the vouchers sold on August 1 and 3 were the same, the sale mechanism through which this item was sold differed between the two days. On August 1, the vouchers were auctioned via ebay's standard incremental auction which is a slightly modified, open-bid second-price auction, starting from a price of 1.00€ .[8] Buyers could therefore influence the final price through their bid, which was also an upper cap on the price they had to pay in case that their bid was the highest. In contrast, vouchers sold

[8] More precisely, in eBays "proxy"-auction, a bidder can submit a bidding cap. Starting from an initial price, eBay than raises the price to the second-highest cap plus an increment as long as this does not exceed the highest cap. The increment depends on the price and is 1.00€ or less for the auctions reported here. This so-determined price is displayed and bidders can re-can raise their bidding cap if they do so before the fixed ending time of the auction. The winner then has to pay the final price. More information on eBay's sales mechanisms can be found in Hasker and Sickles (2010).

on August 3 were offered for a fixed price of 66.00€. In this sale, buyers could not influence the price; it was a take-it-or-leave-it offer. Except for these differences in the sale mechanism, all the procedures and transactions characteristics, e.g. payment options, the seller, the sale platform, shipping procedures, and shipping costs were the same.[9]

eBay allows and encourages its members to mutually review their transactions. Part of such a review is that buyers can rate sellers along several dimensions and give an overall feedback score which is either "negative", "neutral" or "positive". The seller's feedback score which is the sum of positive overall feedbacks minus the sum of negative overall feedback (neutral overall feedback counts zero) is prominently displayed beside a seller's account name. Each seller also has a publicly accessible seller profile. It displays, for a limited time, for each of the reviews which the seller got the following information: The review's overall feedback score, i.e. whether it was positive, neutral or negative, the time when the review was left, the offer to which the review refers, the username of the buyer who left the review, and a short text comment written by the reviewing buyer. For each review, the feedback score of the reviewing customer's account is also displayed next to this customer's account name. Note that in contrast to feedback from buyers for sellers, feedback from sellers for buyers can only be positive or not be given at all, a rule which eBay had previously introduced to prevent that buyers and sellers exchange feedback in a reciprocal manner (see Bolton et al., 2013).[10]

Starting from August 1, I obtained the reviews and associated data for the two offers from the seller's review page for forty consecutive days. This paper looks, for reasons which will later become clear, on data covering multiples of six days after the initial auction data. This means that in the following, I will use data from reviews for the auction which were left in the 36 days from August 1 and, for comparison, data from reviews for the fixed-price sale left from August 3 on over the following 36 days.[11] Table 1 below shows the summary statistics for this data which was obtained from reviews displayed on the seller's profile page. For reviews which refer to the auction the price is, on average, 13.09€ higher than for the fixed price. In fact, just 9 out of 3,575 (0.25%) reviews for the auction were for a transaction which yielded a final price which was below the 66.00€ of the fixed-price sale. The table also displays the time since the auction or fixed sale for when the review was left. It is counted in days, starting from

[9]For both, the auction and the buy-it-now sales, the final sale price was subject to an additional shipping fee of 2.50€. Also, in both cases the voucher came with an additional 10.00€-discount coupon which could be applied for later regular ticket purchases in the webstore of the railway company.

[10]Until May 2008, sellers could give buyers not only positive or no feedback but also negative or neutral feedback.

[11]All crucial results remain unaffected when the additional four days of observations for the auction reviews and the additional two days for the fixed price reviews are used.

	AUCTION		FIXED PRICE	
	mean	(s.d.)	mean	(s.d.)
Price	79.09	(13.64)	66.00	(0.00)
Days since transaction	12.40	(7.47)	13.41	(7.92)
Buyer's Score	207.99	(504.10)	182.79	(628.18)
$\leq$10	11.1%		11.3%	
$\leq$100 (and $>$10)	43.8%		47.0%	
$\leq$1000 (and $>$100)	41.8%		39.3%	
$>$1000	3.3%		2.4%	
Observations	3,575		15,175	

Table 1. Description of the price paid, the days passed between when a review was written and the transaction it refers to, and the feedback score of the buyer who left the review, grouped by reviews for auctions and fixed price transactions. 25 (0.7%) of the buyers who left a review for the auction of and 131 (0.9%) of the buyers who did so for the fixed price hid their feedback scores; Buyer's Score-statistics omit these observations (see footnote 12).

zero in case the review was left on the day of the transaction. This zero date therefore corresponds to August 1 for the auction and August 3 for the fixed-price sale. On average, feedback for the auction is left about one day earlier than for the fixed price.

The variable "Buyer's Score" reflects the feedback score associated with the reviewing customer's account.[12] Thus, it corresponds to the sum of positive feedback previously received by the reviewing customer net of potential previously received negative feedback (see foonote 10). In general, eBay members have a very high rate of positive feedback scores, with an average over 98% and a median rate of 100% (see Bolton et al., 2013; Nosko and Tadelis, 2015). A reviewing buyer's feedback score can therefore be taken as an approximation of this customer's eBay-experience, although it is actually a lower bound on the number of prior transaction in which the reviewing customers has been involved. As these scores are strongly dispersed, I created four categories, defined by whether the score is weakly less than 10 or whether it surpasses the thresholds of 10, 100, or 1000. Although to some degree arbitrary, the first category for a buyer's feedback score can be thought of belonging to relatively fresh eBay-members. The second and third categories then contain experienced and very experienced members. Those in excess of at least 1000 prior transactions are very likely to be professional sellers themselves who acted as buyers in this transaction. Generally speaking, buyers are fairly experienced: For both sales mechanisms, just around 11% of buyers had a feedback score of ten or less while only around 3% of the buyers were supposedly professionals with at least 1000 transactions; the rest lies in between.

Regarding the information which customers had about the two offers, the following can be said: In

[12] eBay allows, but discourages, its member to conceal their feedback score as long as they only act as buyers. Here, this applies to less than 0.8% of the observations. These observations are omitted for all analysis involving the buyer's score.

principle, the whole sequence of offers for each of the ten days was fixed before-hand and advertised on the seller's eBay page. Buyers could, therefore, have known in advance that exactly the same voucher which was initially auctioned on August 1 was also going to be sold for a fixed price two days later. However, the advertisement for the sales campaign (for example via banners on eBay and other web pages) mostly focused on the respective day's special offer and lead directly to the pages of these offers. Thus, if customers did not explore the seller page, they might have easily missed the fact that a fixed-price sale followed the auction when they participated in the auction.

However, after the auction had ended, customers' information regarding the fixed-price offer could change. Successful bidders received a confirmation email which listed the obtained item and the final price, together with further information regarding the payment and shipping procedures. Right below this essential information, the confirmation email also contained a list of the seller's upcoming sales. For those who obtained the voucher in the auction on August 1, this confirmation email therefore also featured a relatively salient advertisement for the fixed-price sale two days later. Also, the accompanying advertisement campaign continued to be focused on the respective day's special offer, including advertisement for the fixed-price sale on August 3. In addition, several news pages started to report on the rather particular sequence of sales mechanisms after the auction had ended. Customers who had obtained a ticket in the auction and were initially not aware of the following fixed-price offer might thus have learned about it only ex-post, after they had won an auction. This could have lead to a shift in their reference price which, according tho the above theory, influenced their feedback. Based on this setting and the data obtained from it, I will explain in the next section more precisely how these changes in reference prices and their effects on feedback will be identified.

4 Predictions and identification strategy

In the following, I will lay out how a change in customer's reference prices can affect subsequent feedback and which predictions this implies in the present setting. For this, I will take a top-down approach by going from predictions which address coarse but salient features of the data to more detailed predictions which explicitly address the channels underlying feedback. Such feedback can be either negative, neutral or positive. Positive feedback is the overwhelming norm on eBay with several studies concluding that any other feedback, including neutral feedback, is considered to be a bad evaluation (see Dellarocas and Wood, 2008; Cabral and Hortaçsu, 2010; Bolton et al., 2013; Cabral and Li, 2015; Nosko and Tadelis, 2015). Therefore, "non-positive", i.e. neutral or negative, feedback is, in the sense of the above model, considered unfavorable towards the seller and will be the dependent variable of interest.[13]

The effect which I examine is how feedback is affected by an ex-post shift in the reference price. In face of previous research which shows that customers take prices they observe or paid themselves as reference prices (see Shafir et al., 1997; Ariely and Simonson, 2003; Simonsohn and Loewenstein, 2006; Amir et al., 2008; Weaver and Frederick, 2012), the auction price is natural starting point for a successful bidder's reference price. However, if customers later become aware of a lower fixed-price offer, for example through advertisement in the confirmation email or banners on the day of the fixed-price sale, such an offer can shift their reference price downwards. If this downward shift and the customer's reference-dependence are large enough, a previously positive assessment can be rendered negative so that non-positive feedback can be issued (see Proposition 1). Since, except for the difference in the sales mechanism, all other objective characteristics of the sale were the same, this leads to the following, first prediction:

Prediction 1. *The share of non-positive feedback in reviews for auctioned vouchers is larger than in reviews for vouchers sold for the fixed price.*

While the item and seller characteristics are constant across the auction and the fixed-price offer, the buyers who left a review can be different across these sales mechanism in way which affects feedback-giving. For example, customers who left feedback for the auctions seem to be slightly more experienced than those in the fixed-price offers (see table 1). Previous research shows that socially motivated behavior such as such as giving feedback is moderated by market participants' experience and their own reputational stakes (List, 2003, 2006) which might lead to selection issues. Also, precautionary buying motives in face of potential capacity constraints might have lead to a selection of buyers with high valuations into the auction (however, this would rather lead to a prediction opposite than the above).

To address such concerns I will derive predictions which are only based on data from the auction reviews. They are, therefore, conditional of a potential selection into this mechanism having already occurred and can not be caused by it. Within this sales mechanism, one can then exploit the fact that the auction prices varied. This variation moderates how observing a lower reference price can lead customers to ex-post negatively re-assess a transaction. This is because a downward shift in the

[13]eBay also has a dedicated help page to explain under which conditions negative and neutral feedback is warranted Thus, it treats *both* these feedback options as non-standard, adverse judgments (see footnote 16). In terms of the model, eBay's feedback system is therefore represented by $F = \{f_{-2}, f_{-1}, f_0, f_1\}$ with successive elements referring to negative/neutral/no/positive feedback, respectively.

assessment is increasing in the distance between the auction price and the lower fixed price (see Corollary 1) and leads to the following prediction:

Prediction 2. *The probability of observing non-positive feedback for an auctioned voucher is increasing in the price in which the auction ended.*

While the above prediction is consistent with an effect of reference price shifts on feedback, another mechanism could, in principle, also explain such an effect. Even if customers had no reference-dependent utility ($\mu = 0$) they could experience a negative shock in experienced utility ($e_t < 0$) after they bought the voucher. If this shock is sufficiently large, it might lead to a negative assessment which can then result in non-positive feedback. The probability that for a given size of such a shock, the negative threshold required for such non-positive feedback is undercut decreases in the magnitude of the, then negative, assessment. This magnitude is increasing in the paid price such that the above prediction could also emerge through a shock in experienced utility, not necessarily a shift in reference prices. It would therefore be ideal to have an indicator of whether a customer's reference point was revised or not in order to directly identify a reference price effect. Unfortunately, this requires to know what clients *perceive* which is generally hard to measure, in particular so for observational field data. However, a variable which indicates whether changes in customers' user experience are more likely relative to changes in their reference price serves essentially the same purpose in identifying the driver of a reference price effect. To obtain such an indicator, I will exploit the specific time structure of the sales campaign and when information regarding the fixed-price offer arrived, relative to other information.

After the auction, successful bidders could learn through several channels about the fixed-price offer. During this time, these customers' information could therefore change and led them to revise their reference prices. In contrast, transaction-specific information regarding their acquisition, e.g. experiences during the associated train ride or whether the voucher was actually sent by the seller, remained constant for a while. This is due to the fact that the shipment of the paper voucher consumed time and was initiated only after the customer's money transfer had arrived at the seller's account. Relative to the information she had at the time of the auction, a successful bidder could therefore not get any new information regarding the idiosyncratic transaction before the voucher had arrived by post. Her reference price, however, might have been affected in this time period through the information she obtained regarding the other fixed-price offer.

More formally and in terms of the above model, this means the following: In period $t = 1$, which is defined to immediately follow the transaction and ends at the earliest day when the article could have arrived, it holds that $Var[r_t] > Var[\epsilon_t] = 0$ (i.e. there is variation in the reference price but not in user experience). This inequality does not need to hold in subsequent periods where variations in buyer experience could outweigh variations in reference prices. Suppose this is true and $Var[r_t]$ is large relative to $Var[\epsilon_t]$ only in $t = 1$ but less large or smaller thereafter. The price effect identified from the auctions' multiple end prices, should then, if caused by a reference point shift, be more pronounced in this initial period than in later periods.

To determine this initial time period during which only reference prices but not the user experience could change, I manually checked the buyer comments for statements which indicate that a voucher has arrived. I found the first such statement in a review for an auction dated on August 7. Reassuringly, the first such comment in a review for the fixed price also appeared on the same day. This suggests that the seller dealt with the after-sale logistics for these identical items in the same way. Therefore, in the first six days from the day of the auction, no vouchers should have reached the successful bidders. This conclusion is confirmed by Figure 1. The grey bars in it display the temporal pattern of when reviews for the auction were left:

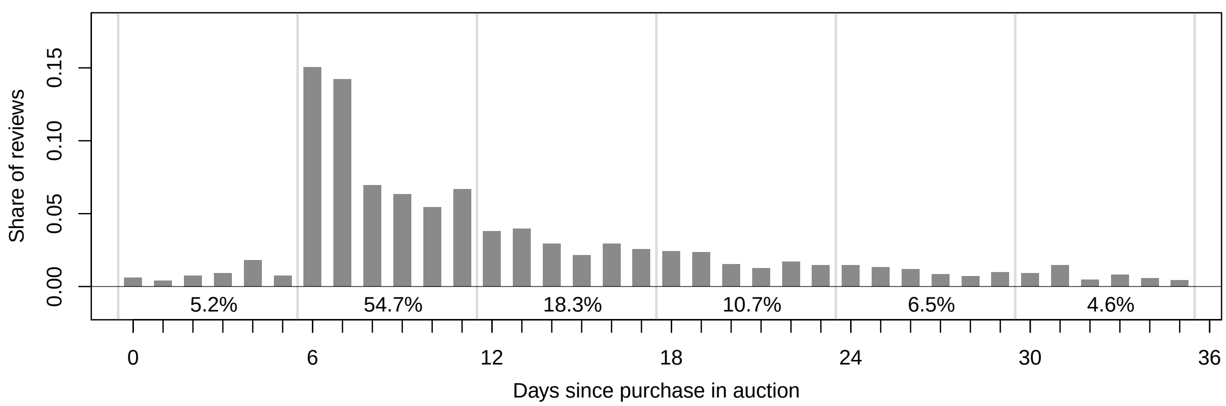

Figure 1. *Grey bars:* Share of total reviews for auctions on a given day. *Percentage numbers:* Share of reviews per bin of six consecutive days (bins separated by vertical lines). Total number of auction reviews: 3,575.

In the first six days from August 1, the day of the auction offer, relatively few reviews occur. During each of these initial six days, less than 2%, together 5.2%, of all the 3,575 reviews for auctions were left. Six days after the auction, on August 7, the daily rate suddenly spikes to almost 15% and remains relatively high for all days in the second six-day-bin which accommodates 52.4% of all reviews for the auction. This is consistent with the notion that, starting from August 7, customers got the voucher and that from this date on, uncertainty regarding the user experience, e.g. the quality of the train ride or whether the voucher would arrive at all, resolved and triggered customers to leave reviews. Thus, from day 6 (August 7) on, new information regarding the customers' idiosyncratic transaction could have affected the their feedback, in addition to the information regarding the fixed-price sale.

The importance of reference point shifts in the first six days from the auction date, relative to later dates is also confirmed by an analysis of the text comments which sellers left with their reviews. I inspected these comments and marked those which refer explicitly to the seller's sales strategy of first selling via an auction and then via a fixed price. As a first evidence regarding the negative feelings this triggered, it is worthwhile to note that most of these comments were written in a hostile and complaining manner. More important regarding the proposed identification strategy is, however, the timing of these comments. The connected black triangles in Figure 2 below display the share of reviews which feature such pricing-related comments relative to the total number of reviews left on a given day. On the day of the auction and the day thereafter (day 0 and 1), no such comments are observed. Then, on the second day after the auction, when the fixed-price sale took place (day 3), the share of daily reviews which contains such comments jumps to more than 36% and stays high for the next three days. A week after the auction, on day 6, when the vouchers started to arrive by post, the share of such comments falls sharply and stays relatively low. This corresponds exactly to the proposed pattern of how the salience of reference price related information relative to other information behaves over time and how this affects customers' perceptions.

The above visual impression regarding the timing of pricing-related comments can also be confirmed statistically. For this, I use the binary variable which indicates a comment referring to the sellers' pricing strategy as the dependent variable in a regression on six dummies for each of the six-day-bins covering the data's 36 days. The estimated coefficients of these six dummies are displayed as horizontal dark grey lines in each the six-day-bins in Figure 2. The light grey rectangles indicate the associated 95%-confidence intervals. The results of this and further regressions which control for additional factors can be found in the Appendix B (table 5). They clearly show that the share of price-related comments, which is around 23% during the first six-day-bin, is significantly lower by 14 to 20 percentage points for each of the following five six-day-bins ($p < 0.01$). Overall, these findings are highly consistent with the notion that for the first six days, the news of the fixed-price offer and the associated reference price shift were a stronger determinant of reviews than in the following days, when customers' feedback was relatively more influenced by customers' experiences related to their idiosyncratic transactions. Accordingly, the

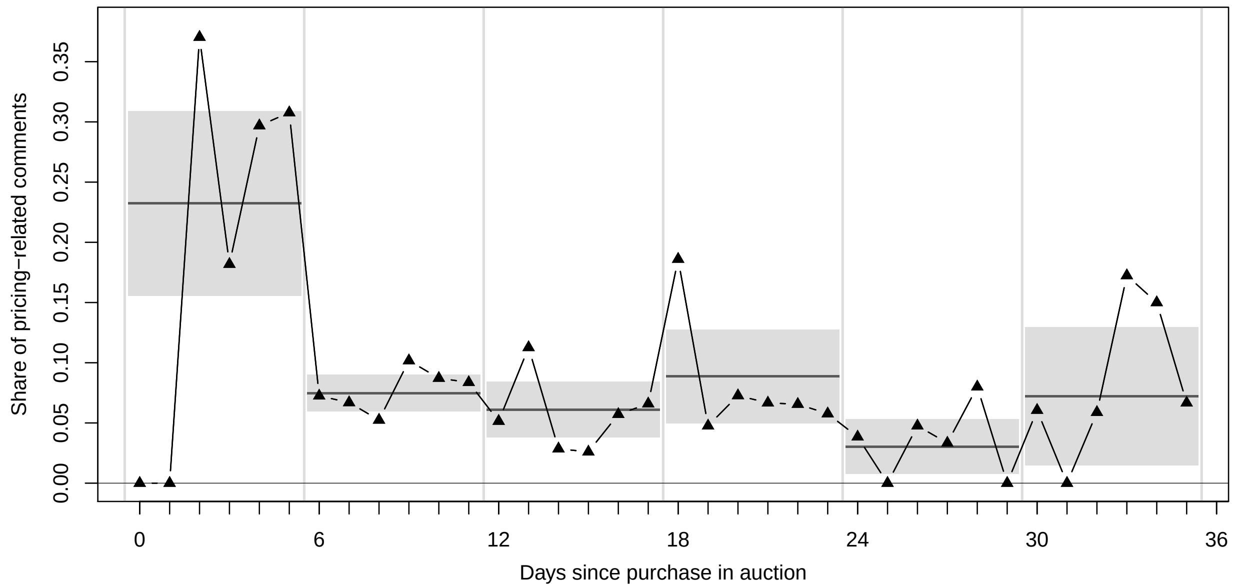

Figure 2. *Black triangles:* Share of reviews on a given day in which the comment refers to the seller's pricing strategy. *Grey rectangles and horizontal lines:* 95%-confidence interval and point estimates of coefficients indicating the probability that such a comment occurs within each bin of six days (separated by vertical lines).

price effect, if it is caused by a shift in the reference price, should be more pronounced in these first six days than later. This leads to the following prediction which explicitly addresses a price effect induced by reference price shifts:

Prediction 3. *The effect of higher auction prices on the probability of non-positive feedback for auctioned vouchers is stronger during the first six days from the auction date, when the reference point shift caused by the fixed-price offer was more salient relative to other information, than in later periods.*

Note that the above reasoning does not predict a gradually decreasing price effect over time. It rather stipulates a sharp decline in the price effect's magnitude after six days and no further change thereafter. Also, note that a test of the above prediction does not rely on any manually coded variable such as the dummy which indicates pricing-related comments. Rather, this manually coded dummy was used to derive this prediction. Its test will be based on "hard" data from reviews in eBay's database which are displayed on the seller profile, i.e. an auction's final price and the date when a review for it was left. Any imprecision or subjective wiggle room in the manual coding of pricing-related comments would thus negatively affect the reasoning which leads to the above prediction. In consequence, such errors would make it harder to confirm the above prediction but do not impose a hazard with respect to a false positive. By the same reasoning, a potential situation where the size of variation in reference prices relative to variations in other information is not only large in the first six days (as assumed in deriving the above prediction) but also in later periods would also increase the potential for a false negative. However, it would not create any problem with regards to a false positive in test of Prediction 3, i.e. in detecting price effects which are induced by reference price shifts. With this in mind, I will now present the results for the tests of this and the other two predictions.

5 Results

I start with a test of Prediction 1, i.e the share of non-positive feedback in reviews for auctions compared to such feedback in reviews for the fixed-price sale. In the auction, 3.4% of all reviews entail a non-positive feedback. This share of is somewhat higher than the usual rate of such feedback on eBay which is, on average, lower than 2% (see Bolton et al., 2013). This difference could be due to the seller- or item-specific effects. However, for the same seller and the same item, the share of non-positive feedback is four times higher for reviews which refer to an auction, at a level of 13.6%. To examine this in more detail and to control for differences in the reviewing customers' experience, I estimate the following regression model:

$$\mathbb{1}[f_i < 0] = \alpha + \beta \cdot Auction_i + \sum_{s=1}^{3} \delta_s \cdot \mathbb{1}[Buyer's\,Score_i > 10^s] + \epsilon_i \tag{3}$$

The dependent variable in the above is an indicator which takes a value of one if the statement in the brackets is true, i.e. if the feedback in review i is non-positive. $Auction_i$ is a dummy which is equal to one if this review was for an auction. With respect to tests of Prediction 1, it is the main variable of interest. The three indicators for whether $Buyer's\,Score_i$ surpasses the respective power-of-ten-threshold (but not the next-highest) correspond to the lower bound on the number of the reviewing buyer's hitherto transactions and therefore measure her experience. Unobserved idiosyncratic features of the transaction which affect ratings, for example user experience, are captured by the error term ϵ_i. In the main text, I will always present the results of linear probability models for regressions of the above type. This is done for easier exposure and interpretation, in particular with regards to later-used models which feature interaction-terms and which are somewhat more difficult to interpret when non-linear models such as Probit are used. However, the results of such non-linear models, which are collected in Appendix C, do not differ in any meaningful manner from those presented here.

	(1)	(2)	(3)	(4)
Auction	0.102***	0.102***	0.100***	0.100***
	(0.008)	(0.008)	(0.007)	(0.007)
Buyer's Score >10		-0.018**		-0.015**
		(0.008)		(0.006)
Buyer's Score >100		-0.015**		-0.012*
		(0.008)		(0.006)
Buyer's Score >1000		-0.003		0.004
		(0.016)		(0.014)
Constant	0.034***	0.048***	0.032***	0.043***
	(0.002)	(0.007)	(0.002)	(0.005)
First reviews only	no	no	yes	yes
Observations	18,750	18,594	15,857	15,721
Adjusted R^2	0.032	0.032	0.028	0.028

Table 2. OLS estimates of regressing an indicator for non-positive feedback on a dummy for whether the review was for the auction and dummies for the reviewing buyers' own feedback score. Columns 1 & 2 use all collected reviews and report standard errors clustered on the buyer level. Columns 3 & 4 only use the first review a buyer posted and report robust standard errors. Significance levels: *** p<0.01, ** p<0.05, * p<0.1

Table 2 above presents the OLS-results for regression model (3). The estimates in the first column reflect the previously indicated difference of 10.2 percentage points between the sales mechanisms: The share of non-positive feedback is significantly higher in reviews for auction sales than in reviews for fixed-price sales. The results in column 2 show that adding controls for the buyer experience via their own feedback scores leaves this difference unchanged. Thus, this result is not caused by differences in buyer experience. Taken for themselves, these controls indicate that intermediately experienced buyers are significantly less likely to give non-positive feedback than inexperienced buyers, the omitted category. In contrast, buyers with at least 1000 prior transactions do not differ significantly from them in their propensity to give non-positive feedback. However, the magnitude of these pure experience effects is relatively small, compared to the adverse effect of selling via the auction.

A potential confounder for this results is that some customers bought multiple vouchers. This allowed them to issue multiple reviews and this may have disproportionately affected reviews for one of the sales mechanisms. I therefore repeated the analysis when only the first review which each distinct buyer left is used for the analysis. Column 3 and 4 in table 2 present the corresponding results. They do not differ in any meaningful way from the previous results when the whole sample is used. These findings are in support of Prediction 1 and can be summarized as follows:

Result 1. *The share of non-positive feedback in reviews for vouchers obtained in the auction is significantly higher in the auction as compared to reviews for vouchers obtained for a fixed price.*

The above shows that transactions for the same exchanged item, the same seller, and on the same platform receive much more unfavorable feedback when they were based on an auction as opposed to the subsequent fixed-price sale. It is suggestive to attribute this to the fact that almost all auctions for which a review was left resulted in a higher price than the fixed price of 66€. To analyze this price effect, I exploit the price variations in the auctions and estimate a regression model similar to (3). The only difference is that instead of the $Auction_i$-dummy, $\Delta_{10}Price_i$ is included as the key independent variable. This variable measures the effect of the price paid in the auction on the probability of getting non-positive feedback. To make the interpretation of its coefficient easier, it is not the absolute price paid in the auction but the difference to the fixed price of 66€, divided by 10. Thus, the point estimate refers to the change in the probability of leaving a non-positive feedback, associated with a change in the price difference by 10€. Table 3 reports the OLS-results for this model.

	(1)	(2)	(3)	(4)
Δ_{10}Price	0.023***	0.023***	0.024***	0.024***
	(0.006)	(0.006)	(0.006)	(0.006)
Buyer's Score $>$10		-0.036		-0.035
		(0.029)		(0.025)
Buyer's Score $>$100		-0.044		-0.038
		(0.029)		(0.025)
Buyer's Score $>$1000		0.010		0.034
		(0.051)		(0.050)
Constant	0.105***	0.139***	0.101***	0.131***
	(0.010)	(0.027)	(0.009)	(0.023)
Observations	3,575	3,550	2,283	2,265
First reviews only	no	no	yes	yes
Adjusted R^2	0.008	0.009	0.009	0.010

Table 3. OLS estimates of regressing an indicator for non-positive feedback on the difference between the auction price and the fixed price divided by 10 and dummies for the reviewing buyers' own feedback score. Columns 1 & 2 use all collected reviews and report standard errors clustered on the buyer level. Columns 3 & 4 only use the first review a buyer posted and report robust standard errors. Significance levels: *** $p<0.01$, ** $p<0.05$, * $p<0.1$

The results show a clear price effect: The higher the price paid in the auction, the more likely is that a review left entails a non-positive feedback for the transaction. On average, for each 10€ paid more in the auction than for the fixed-price offer, the probability of non-positive feedback rises by about 2.3 percentage points, independently of whether controls are added or whether multiple reviews from the same customer are used. Therefore, the following finding, in line with Prediction 2, can be stated:

Result 2. *The share of non-positive feedback in reviews for vouchers obtained in the auction increases significantly in the price paid.*

The above findings are consistent with a reference price shift among those who obtained the voucher in the auction. However, it represents the average price effect across the whole sample period. In the following, I will explicitly link this price effect to a reference price shift by decomposing it along the temporal structure of the post-sale procedures. For this, I will exploit that during the first six days from the auction date until the voucher started to arrive by post, no new information regarding the specific transaction a customer was involved occurred. Knowledge about the fixed price could, however, change within these six days so that reference prices could be updated in this period. In consequence, the price effect, if it is caused by a shift in customers' reference points, should be more pronounced during the initial period than later. To test this, I estimated a regression model where the price effect is measured separately for the initial six days and the sample's five remaining six-day-bins. That is, the following regression model is estimated:

$$\begin{aligned}\mathbb{1}[f_i < 0] = \alpha + \beta_1 \cdot \Delta_{10} Price_i + \sum_{t=2}^{6} \beta_t \cdot \Delta_{10} Price_i \times SixDayBin\#t_i \\ + \sum_{t=2}^{6} \gamma_t \cdot SixDayBin\#t_i + \sum_{s=1}^{3} \delta_s \cdot \mathbb{1}[Buyer's\ Score_i > 10^s] + \epsilon_i\end{aligned} \tag{4}$$

The dependent variable is, as in the preceding regressions, whether a review entails non-positive feedback. Also as before, it features indicators for the reviewing customer's own feedback scores as control variables and the price difference to the fixed-price offer. The regression model also includes five dummies denoted by $SixDayBin\#t_i$. Their values indicate during which bin of six consecutive days, starting from the auction date, a review was written with the first six-day-bin being the omitted category. The coefficient for the price difference and its interaction with these $SixDayBin\#t_i$-indicators are the main independent variables of interest. As the first six-day-bin is the omitted category, the coefficient β_1 measures the price slope in this period. The interacted price-difference-coefficients (β_2 through β_6) explicitly capture differences of the price effect in later periods whereas the coefficients for the non-interacted six-day-bins (γ_2 through γ_6) capture differences in feedback relative to the first six days when the price difference is zero. Prediction 3 than states that the price slope during the first six days, given by $\hat{\beta}_1$, should be more pronounced than the price slope in a later period $t \geq 2$, given by $\hat{\beta}_1 + \hat{\beta}_t$. Thus, the estimates $\hat{\beta}_t$ for the interaction terms are predicted to be negative.

Table 4 reports the results when the above model is estimated by OLS for the reviews which refer to the auction. The results are in line with the hypothesis. The non-interacted price effect is strong

	(1)	(2)	(3)	(4)
Δ_{10}Price	0.137***	0.137***	0.127***	0.126***
	(0.028)	(0.028)	(0.031)	(0.031)
Δ_{10}Price$\times$SixDayBin#2	-0.129***	-0.128***	-0.119***	-0.118***
	(0.028)	(0.029)	(0.032)	(0.032)
	[0.206]	[0.174]	[0.210]	[0.179]
Δ_{10}Price$\times$SixDayBin#3	-0.129***	-0.131***	-0.121***	-0.119***
	(0.029)	(0.029)	(0.033)	(0.033)
	[0.375]	[0.453]	[0.544]	[0.482]
Δ_{10}Price$\times$SixDayBin#4	-0.116***	-0.116***	-0.105***	-0.105***
	(0.035)	(0.035)	(0.037)	(0.037)
	[0.322]	[0.319]	[0.276]	[0.280]
Δ_{10}Price$\times$SixDayBin#5	-0.142***	-0.139***	-0.153***	-0.151***
	(0.031)	(0.032)	(0.034)	(0.034)
	[0.772]	[0.897]	[0.043]**	[0.062]*
Δ_{10}Price$\times$SixDayBin#6	-0.101**	-0.109***	-0.080*	-0.087**
	(0.041)	(0.040)	(0.043)	(0.043)
	[0.235]	[0.334]	[0.131]	[0.195]
SixDayBin#2	-0.102*	-0.100*	-0.127**	-0.124**
	(0.058)	(0.058)	(0.060)	(0.060)
SixDayBin#3	-0.112*	-0.110*	-0.125**	-0.125**
	(0.059)	(0.059)	(0.062)	(0.062)
SixDayBin#4	-0.074	-0.071	-0.074	-0.070
	(0.066)	(0.066)	(0.068)	(0.068)
SixDayBin#5	-0.123*	-0.124*	-0.099	-0.099
	(0.064)	(0.064)	(0.069)	(0.069)
SixDayBin#6	-0.072	-0.065	-0.139*	-0.127*
	(0.079)	(0.079)	(0.071)	(0.071)
Buyer's Score $>$10		-0.031		-0.030
		(0.028)		(0.025)
Buyer's Score $>$100		-0.047*		-0.038
		(0.028)		(0.025)
Buyer's Score $>$1000		0.000		0.028
		(0.052)		(0.051)
Constant	0.206***	0.237***	0.222***	0.247***
	(0.056)	(0.062)	(0.059)	(0.063)
First reviews only	no	no	yes	yes
Observations	3,575	3,550	2,283	2,265
Adjusted R^2	0.069	0.069	0.077	0.076

Table 4. OLS estimates of regressing an indicator for non-positive feedback on the difference between the auction price and the fixed price divided by 10, a dummy for the six-day-bin since the auction in which the review was written, its interaction with the price variable, and dummies for the reviewing buyers' own feedback score. Columns 1 & 2 use all collected reviews and report standard errors clustered on the buyer level. Columns 3 & 4 only use the first review a buyer posted and report robust standard errors. Numbers in square-brackets refer to the p-value of a F-tests on the null that the coefficient on Δ_{10}Price plus the coefficient on the respective interaction term Δ_{10}Price$\times$SixDayBin#t are equal to zero. Significance levels: *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

and significant. It corresponds to an increase of around 13 percentage points in the probability of non-positive feedback for each 10€ a customer paid above the fixed price if the review was left in the first six days from the day of the auction. This effect is much stronger, by a factor larger than five, than the average price effect which was previously estimated over the whole 36 days (see table 3 above). It is also much larger than for all reviews left within one of the later six-day-bins. The coefficients on the corresponding interaction terms are consistently estimated to be significantly negative. In fact, the price slope for the later periods, given by the coefficient on the un-interacted plus the respective interacted price difference, often indicate a price effect close to zero in these periods. Formal tests of this, i.e. F-tests on the null that $\hat{\beta}_1 + \hat{\beta}_t = 0$, are in line with this impression. The associated p-values are portrayed in the square-bracket beneath the estimated coefficients for the interaction terms in table 4. In all of these twenty tests, except for two where the implied price slope is slightly negative, the null of a zero price effect cannot be rejected at usual significance levels.

Note that the observed pattern of the price-effect does not indicate a decreasing time effect over time. It rather shows that after the first six days, the price-slope sharply decreases close to zero and stays roughly constant at this level over the sample's remaining days. This is inconsistent with a notion that negative emotions "cool down" over time. This would predict a gradually decreasing price-slope, i.e. more negative coefficients for higher-numbered interaction terms instead of a single jump after the initial six days. This is, however, not observed.[14] The findings are rather consistent with the notion that for the first six days, the new information regarding the fixed price was the driving force behind the unfavorable feedback. As soon as other, transaction-related information became salient, this reference price effect ceased to determine feedback.[15] Therefore, the following result, which directly confirms Prediction 3's suggested timing of price effects as caused by reference price shifts, can be stated:

Result 3. *The significant positive effect of the price paid on the probability of observing non-positive feedback is concentrated in the first six days from the day of the auction when reference prices were more salient than later. After these initial six days, the price effect becomes essentially zero and insignificant.*

[14] I also estimated a regression model in which the days since the auction form a time trend which is interacted with the price variable. If this model is estimated for a period covering the first five six-day-bins, this interaction term is significantly negative, supposingly because it captures the first six days' strong effect on the price slope. I then estimated the same model with data from the last five six-day-bins which have almost the same number of observations, and therefore potential power, as the first five bins. However, the interacted time trend and the price effect in this second regression are estimated at almost exactly zero (and are not significantly different from it). Thus, the effect is not gradually decreasing.

[15] This is also confirmed by simple regression with an indicator for non-positive feedback as the dependent and the auction price and a constant as the sole independent variables. When estimated with auction data from only the first six-day-bin one gets an R^2 of 0.228. When the same model is estimated with data from all the other six-day-bins, the R^2 decreases to 0.015. Outside the first six days, feedback is therefore much better explained by unobserved factors other than the price.

6 Discussion & Conclusion

In this paper, I demonstrate how ex-post reference price shifts can adversely affect customer behavior in an auction context. This manifests through unfavorable public feedback which successful bidders are more likely to leave the higher the auction price is, even though their own bids bound this price. I show that this price effect is due to successful bidders who learn, after they have won the auction, about an almost identical offer which only differed in the fact that it was sold for a lower, fixed price. This information is not relevant for assessing the objective value of the idiosyncratic transactions in which these auction customers had participated in. However, it can shift their reference prices downward and thereby negatively affect how they assess their transaction. This then causes, moderated by their social preferences, a negative effect on the seller's reputation.

Before I discuss the consequences of these findings, recall that they are based on reviews which were left voluntarily by customers. For them, the outcome of the transaction they participated in and their motivation to give feedback was important enough that they considered it worthwhile to give feedback which reflects their assessment. As having obtained the object is a necessary but not a sufficient condition for giving feedback, the point estimates presented here should be taken with some caution and as representative of those customers, not the entire population. However, the main results regarding the difference in feedback across sales mechanism, the price effect, and how it is moderated by the timing of customers' information are all based on *relative* effect sizes within the sample of feedback-giving customers. The point estimates are therefore less important. Also, only these feedback-giving customers' assessments materialize into concrete, observable actions which determine the seller's feedback score. As such feedback scores and ratings crucially affects sales records (see Melnik and Alm, 2002; Livingston, 2005; Houser and Wooders, 2006; Resnick et al., 2006; Anderson and Magruder, 2012) this study's findings have several implications:

First, this paper documents an unintended consequence of pricing patterns which are not uncommon in online markets. Einav et al. (2015) show that auction and fixed-price offers for the same retail goods are often available within close temporal succession. However, auction customers often end up paying more than in the competing fixed-price offers available to them (Ariely and Simonson, 2003; Malmendier and Lee, 2011). Similar pattern can also occur when reverse auctions are used alongside fixed-price offers to sell unused capacities, for example in the travel and hotel industry (the most prominent example being Priceline, see Wang et al., 2009; Gardner, 2012). Sequential sales where auctions preceed fixed-price sales might, at first sight, also seem appealing to sellers in various other situations. For example, first selling via an auction can help a monopolist to construct a demand curve from the observed bids. Based on this, it can then compute a profit-maximizing price for subsequent fixed-price sales. The same sequential sales strategy, though less motivated by profit-seeking concerns, can also be employed to prevent (ticket-)scalping, i.e. the higher-priced resale of rationed goods (Courty, 2003; Roth, 2007; Leslie and Sorensen, 2014). First selling via an auction and then selling a potential remainder for a lower, fixed price reverts and destroys the business model of scalpers. I show that the use of such pricing strategies comes with the caveat of potentially causing antagonism among those customers who have the highest valuation for a seller's product.

The second implication relates to the first. Sellers can often exploit inattentive and under-searching customers, for example by deliberately using overly complicated pricing rules (Grubb, 2015), preventing price comparisons (Ellison and Fisher Ellison, 2009) or shrouding fees (Brown et al., 2010). This paper's results show that customers punish the seller when they realize that they missed a better deal. As long as the cost of such antagonism is not factored into the seller's trade-off regarding the use of obfuscating sales strategies they can, eventually, backfire.

Third, the adverse reactions by customers which are documented here do not only matter for sellers. They can also be relevant for sales platforms which offer different pricing mechanisms. In a recent paper, Einav et al. (2016) show that the share of auctions on eBay has decreased from about 65% in summer 2009 to just over 15% four years later. In line with the findings presented here, they attribute this fact to the costs of researching the available offers and which are necessary to make a good deal in an auction. This paper's results indicate that, in addition to under-searching customers' direct costs of missing better offers, they experience additional dis-utility when their (supposedly) happy feelings about a successful bid reverse upon later discovering a lower-priced fixed-price offer. Such additional dis-utility than adds another reason for customers' distaste for auctions which is then reflected by sellers' declining use of this sales format. This is reinforced by this paper's finding that sellers are also punished *directly* for using auctions in the presence of competing fixed-price offers.

Fourth and finally, the results presented here also have relevance for the functioning of reputation systems more generally. They are often seen as a way to give market participants some reputational leverage in order to prevent fraudulent behavior and to ensure sufficient quality. This is particular relevant in rather anonymous online situations where traditional mouth-to-mouth reputation cannot fulfill this rule (see Dellarocas, 2003; Bolton et al., 2004; Jin and Kato, 2006; Cabral and Hortaçsu, 2010; for a recent review of the topic see Tadelis, 2016). I show how a reference price shift, i.e. a purely psychological process, leads customers to rate a seller unfavorably. Such feedback does therefore not reflect objectively negative elements of an idiosyncratic transaction such as a delayed shipment or a faulty item. It rather represents the customer's subjective negative experience which was caused by ex-post observing another offer. However, once feelings about a seller's overall sales sequence rather than objective and transaction-specific facts determine feedback, its informational value in the latter dimension diminishes.[16] In consequence, reviews from antagonizing customers can, similar to the problems created by omitted reviews (Dellarocas and Wood, 2008; Nosko and Tadelis, 2015) or fake reviews (Anderson and Simester, 2014; Mayzlin et al., 2014; Luca and Zervas, 2016), create negative externalities from single transactions on the overall quality and informativeness of a market platform's reputation system.

[16] In fact, eBay's user guidelines have tried to prevent this by making it clear that "*Feedback becomes a permanent part of the information about a seller*" and that before buyers give a "*neutral or negative feedback, they should contact the seller and try to resolve problems*". Importantly, these rules also state that feedback should be "*fair and objective*" (translated from eBay.de's feedback rules, retrieved from http://pages.ebay.de/help/feedback/howitworks.html at 09.02.2009).

Appendix A: Proof of Proposition 1

First note that the unrestricted optimum over $\mathbb{R}$, given by $\tilde{f}_t^* \equiv \arg\max_{f_t \in \mathbb{R}} U_t\big(f_n|\psi, A_t\big)$ with $A_t = A_t(e_t, r_t, v, p)$ has to solve $A_t\psi = c'(f_t^*)$. Strict convexity of $c \in \mathcal{C}^2$ with a minimum at zero ensures that this is the only optimum and that c' is strictly increasing. Therefore, c is invertible and $\tilde{f}_t^* = c'^{-1}(A_t\psi)$ holds. Also, as $\tilde{f}_t^*$ is a maximum, $U''(\tilde{f}_t^*|\psi, A_t) = -c''(\tilde{f}_t^*) < 0 = U'(\tilde{f}_t^*|\psi, A_t)$ applies. In consequence, $U'(\tilde{f}|\psi, A_t) \gtreqless 0$ holds for any $\tilde{f} \in \mathbb{R}$ such that $\tilde{f} \lesseqgtr \tilde{f}_t^*$. This means that for some $f_n \in F \subset \mathbb{R}$ the loss $U(f_n|\psi, A_t) - U(\tilde{f}_t^*|\psi, A_t) < 0$ is strictly increasing in $|f_n - \tilde{f}_t^*|$. The restricted solution over F, given by $f_t^* = \arg\max_{f_t \in F} U_t\big(f_n|\psi, A_t\big)$, is therefore uniquely defined and one of the two elements in F which are closest to the unrestricted solution $\tilde{f}_t^*$. For $f_t^* \leq f_n$ to apply, it therefore has to hold that $\tilde{f_t}^* \leq \lambda_n \cdot f_n + (1 - \lambda_n) \cdot f_{n+1}$ where the value of $\lambda_n \in (0, 1)$ is determined by the specific cost function c and the distance $f_{n+1} - f_n$. For any $f_n < \max\{F\}$, the following then holds:

$$
\begin{aligned}
\Pr[f_t^* \leq f_n] &= \Pr\left[\tilde{f}_t^* \leq \lambda_n \cdot f_n + (1 - \lambda_n) \cdot f_{n+1}\right] \\
&= \Pr\left[c'^{-1}(A_t\psi) \leq \lambda_n \cdot f_n + (1 - \lambda_n) \cdot f_{n+1}\right] \\
&= \begin{cases}
\Psi\left(\frac{c'(\lambda_n \cdot f_n + (1-\lambda_n)\cdot f_{n+1})}{A_t}\right) & \text{if } A_t > 0 \\
1 & \text{if } A_t = 0 \text{ and } f_n \geq 0 \\
0 & \text{if } A_t = 0 \text{ and } f_n < 0 \\
1 - \Psi\left(\frac{c'(\lambda_n \cdot f_n + (1-\lambda_n)\cdot f_{n+1})}{A_t}\right) & \text{if } A_t < 0
\end{cases}
\end{aligned}
$$

The proposition is then a direct consequence from the above and the assumptions on Ψ. □

Appendix B: Timing of price-related comments

To verify that comments which related to the seller's pricing strategy were more often issued during the first six days from the day of the auction on, compared to later periods, the following regression model was estimated, using data from reviews for the auction:

$$\mathbb{1}[CommentPricing_i] = \alpha + \sum_{t=2}^{6} \beta_t \cdot SixDayBin\#t_i + \sum_{s=1}^{3} \gamma_s \cdot \mathbb{1}[Buyer's\ Score_i > 10^s] + \epsilon_i \quad (5)$$

The dependent variable is a manually coded dummy which indicates whether a comment in the auction referred to the pricing strategy of the seller (selling first via an auction and then via a fixed price). The series of $SixDayBin\#t_i$-dummies indicate in which of the respective mutiples of six days, starting from the day of the auction, the respective review was written. This is the same variable as used in regression model (4); the first six days are therefore the omitted category. Table 5 reports the results when the above model is estimated by OLS using either all observations from auctions or only each buyer's first review and when controls for the buyer's feedback score are successively added. Results remain unchanged when the model is estimated by Probit (see table 6).

	(1)	(2)	(3)	(4)
SixDayBin#2	-0.158***	-0.153***	-0.143***	-0.137***
	(0.040)	(0.040)	(0.036)	(0.036)
SixDayBin#3	-0.171***	-0.166***	-0.149***	-0.143***
	(0.041)	(0.041)	(0.038)	(0.037)
SixDayBin#4	-0.144***	-0.139***	-0.119***	-0.113***
	(0.044)	(0.044)	(0.041)	(0.041)
SixDayBin#5	-0.202***	-0.197***	-0.179***	-0.173***
	(0.041)	(0.041)	(0.039)	(0.039)
SixDayBin#6	-0.160***	-0.150***	-0.161***	-0.148***
	(0.049)	(0.049)	(0.042)	(0.042)
Buyer's Score >10		-0.024		-0.017
		(0.022)		(0.020)
Buyer's Score >100		-0.035		-0.025
		(0.021)		(0.020)
Buyer's Score >1000		0.062		0.083*
		(0.047)		(0.047)
Constant	0.232***	0.250***	0.216***	0.225***
	(0.039)	(0.044)	(0.036)	(0.039)
First reviews only	no	no	yes	yes
Observations	3,575	3,550	2,283	2,265
Adjusted R^2	0.020	0.024	0.018	0.022

Table 5. OLS estimates of regressing an indicator for a comment in the review which refer to the seller's sales strategy on dummies for the six-day-bin when the review was written and the buyers' own feedback score. Columns 1 & 2 use all collected reviews for the auction and report standard errors clustered on the buyer level. Columns 3 & 4 only use the first review a buyer posted for the auction and report robust standard errors. Significance levels: *** p<0.01, ** p<0.05, * p<0.1

Note: The point estimates and confidence intervals displayed in figure 2 corresponds to the implied estimates from the first column in Table 5 above. The only difference is that they are based on an equivalent regression in which no constant but a dummy for $SixDayBin\#1_i$ is included.

References

Ahmad, H. F. (2015). Endogenous price expectations as reference points in auctions. *Journal of Economic Behavior and Organization 112*, 46–63.

Amir, O., D. Ariely, and Z. Carmon (2008). The Dissociation Between Monetary Assessment and Predicted Utility. *Marketing Science 27*(6), 1055–1064.

Anderson, E. T. and D. I. Simester (2008). Does Demand Fall When Customers Perceive That Prices Are Unfair? The Case of Premium Pricing for Large Sizes. *Marketing Science 27*(3), 492–500.

Anderson, E. T. and D. I. Simester (2010). Price Stickiness and Customer Antagonism. *Quarterly Journal of Economics 125*(2), 729–765.

Anderson, E. T. and D. I. Simester (2014). Reviews Without a Purchase: Low Ratings, Loyal Customers, and Deception. *Journal of Marketing Research (JMR) 51*(3), 249–269.

Anderson, M. and J. Magruder (2012). Learning from the Cloud: Regression Discontinuity Estimates of the Effects of an Online Review Database. *Economic Journal 122*(563), 957–989.

Ariely, D. and I. Simonson (2003). Value Assessment and Decision Dynamics in Online Auctions. *Journal of Consumer Psychology 13*, 113–123.

Avery, C., P. Resnick, and R. Zeckhauser (1999). The Market for Evaluations. *American Economic Review 89*(3), 564–584.

Bartling, B., L. Brandes, and D. Schunk (2015). Expectations as Reference Points: Field Evidence from Professional Soccer. *Management Science 61*(11), 2646–2661.

Battigalli, P., M. Dufwenberg, and A. Smith (2015). Frustration and Anger in Games. *mimeo*.

Bolton, G., B. Greiner, and A. Ockenfels (2013). Engineering Trust: Reciprocity in the Production of Reputation Information. *Management Science 99*(2), 265–285.

Bolton, G., E. Katok, and A. Ockenfels (2004). How Effective Are Electronic Reputation Mechanisms? An Experimental Investigation. *Management Science 50*(11), 1587–1602.

Bordalo, P., N. Gennaioli, and A. Shleifer (2017). Memory, Attention and Choice. *NBER working paper 23256*.

Brown, J., T. Hossain, and J. Morgan (2010). Shrouded Attributes and Information Suppression: Evidence from the Field. *Quarterly Journal of Economics 125*(2), 859–876.

Brynjolfsson, E., Y. J. Hu, and D. Simester (2011). Goodbye Pareto Principle, Hello Long Tail: The Effect of Search Costs on the Concentration of Product Sales. *Management Science 57*(8), 1373–1386.

Cabral, L. and A. Hortaçsu (2010). The dynamics of seller reputation: Evidence from eBay. *Journal of Industrial Economics 58*(1), 54–78.

Cabral, L. and L. I. Li (2015). A Dollar for Your Thoughts: Feedback-Conditional Rebates on eBay. *Management Science 61*(9), 2052–2063.

Card, D. and G. B. Dahl (2011). Family violence and football: The effect of unexpected emotional cues on violent behavior. *Quarterly Journal of Economics 126*(1), 103–143.

CEA (2015). Big Data and Differential Pricing. *Report by the White House's Council of Economic Advisers*.

Chandran, S. and V. G. Morwitz (2005). Effects of Participative Pricing on Consumers' Cognitions and Actions: A Goal Theoretic Perspective. *Journal of Consumer Research 32*(2), 249–259.

Courty, P. (2003). Some Economics of Ticket Resale. *The Journal of Economic Perspectives 17*(2), 85–97.

Dellarocas, C. (2003). The Digitization of Word of Mouth: Promise and Challenges of Online Feedback Mechanisms. *Management Science 49*(10), 1407–1424.

Dellarocas, C. and C. a. Wood (2008). The Sound of Silence in Online Feedback: Estimating Trading Risks in the Presence of Reporting Bias. *Management Science 54*(3), 460–476.

DiTella, R. and J. Dubra (2014). Anger and Regulation. *Scandinavian Journal of Economics 116*(3), 734–765.

Einav, L., C. Farronato, and J. Levin (2016). Auctions versus Posted Prices in Online Markets. *Journal of Political Economy*, forthcoming.

Einav, L., T. Kuchler, J. Levin, and N. Sundaresan (2015). Assessing sale strategies in online markets using matched listings. *American Economic Journal: Microeconomics 7*(2), 215–247.

Ellison, G. and S. Fisher Ellison (2009). Search, Obfuscation, and Price Elasticities on the Internet. *77*(2), 427–452.

Englmaier, F., A. Schmoeller, and T. Stowasser (2016). Price Discontinuities in an Online Used Car Market. *Management Science*, forthcoming.

Ericson, K. M. and A. Fuster (2011). Expectations as endowments: evidence on reference-dependent preferences from exchange and valuation experiments. *Quarterly Journal of Economics 126*(4), 1879–1907.

Fehr, E., O. Hart, and C. Zehnder (2011). Contracts as Reference Points - Experimental Evidence. *American Economic Review 101*(April), 493–525.

Frey, B. S. and W. W. Pommerehne (1993). On the fairness of pricing An empirical survey among the general population. *Journal of Economic Behavior and Organization 20*, 295–307.

Gardner, T. (2012). Priceline introduces Express Deals for hotels. *Los Angeles Times 27.6.2012*.

Gill, B. D. and V. Prowse (2012). A Structural Analysis of Disappointment Aversion in a Real Effort Competition. *American Economic Review 102*(1), 469–503.

Grubb, M. D. (2015). Failing to Choose the Best Price: Theory, Evidence, and Policy. *Review of Industrial Organization 47*(3), 303–340.

Hart, O. and J. Moore (2008). Contracts As Reference Points. *Quarterly Journal of Economics 123*(1), 1–48.

Hasker, K. and R. Sickles (2010). eBay in the Economic Literature: Analysis of an Auction Marketplace. *Review of Industrial Organization 37*(1), 3–42.

Heidhues, P. and B. Köszegi (2008). Competition and Price Variation When Consumers Are Loss Averse. *American Economic Review 98*(4), 1245–1268.

Heidhues, P. and B. Köszegi (2014). Regular prices and sales. *Theoretical Economics 9*(1), 217–251.

Herz, H. and D. Taubinsky (2017). What Makes a Price Fair ? An Experimental Study of Transaction Experience and Endogenous Fairness Views. *Journal of the European Economic Association*, forthcoming.

Hinz, O., I.-H. Hann, and M. Spann (2011). Price Discrimination in E-Commerce? An Examination of Dynamic Pricing in Name-Your-Own Price Markets. *MIS Quarterly 35*(1), 81–98.

Hossain, T. and J. Morgan (2006). ...Plus Shipping and Handling: Revenue (Non) Equivalence in Field Experiments on eBay. *B.E. Journals in Economic Analysis and Policy: Advances in Economic Analysis and Policy 6*(2), 1–27.

Houser, D. and J. Wooders (2006). Reputation in auctions : Theory, and evidence from eBay. *Journal of Economics & Management Strategy 15*(2), 353–369.

Jin, G. Z. and A. Kato (2006). Price, Quality and Reputation: Evidence from an Online Field Experiment. *RAND Journal of Economics 37*(4), 983–1005.

Kahneman, B. D., J. L. Knetsch, and R. Thaler (1986). Fairness as a Constraint on Proft Seeking: Entitlement in the Market. *American Economic Review 76*(4), 728–741.

Kamins, M. A., X. Drèze, and V. S. Folkes (2004). Effects of Seller-Supplied Prices on Buyers' Product Evaluations: Reference Prices in an Internet Auction Context. *Journal of Consumer Research 30*(4), 622–628.

Köszegi, B. and M. Rabin (2006). A Model Of Reference-Dependent Preferences. *Quarterly Journal of Economics 121*(4), 1133–1165.

Lacetera, N., D. G. Pope, and J. R. Sydnor (2012). Heuristic Thinking and Limited Attention in the Car Market. *The American Economic Review 102*(5), 2206–2236.

Lange, A. and A. Ratan (2010). Multi-dimensional reference-dependent preferences in sealed-bid auctions How (most) laboratory experiments differ from the field. *Games and Economic Behavior 68*(2), 634–645.

Leibbrandt, A. (2016). Behavioral Constraints on Pricing: Experimental Evidence on Price Discrimination and Customer Antagonism. *CESifo Working Papers 6214*.

Leslie, P. and A. Sorensen (2014). Resale and rent-seeking: An application to ticket markets. *Review of Economic Studies 81*(1), 266–300.

List, J. (2003). Does Market Experience Eliminate Market Anomalies? *Quarterly Journal of Economics 118*(1), 41–71.

List, J. A. (2006). The Behavioralist Meets the Market: Measuring Social Preferences and Reputation Effects in Actual Transactions. *Journal of Political Economy 114*(1), 1–37.

Livingston, J. A. (2005). How Valuable Is a Good Reputation? A Sample Selection Model of Internet Auctions. *Review of Economics and Statistics 87*(3), 453–465.

Luca, M. and G. Zervas (2016). Fake It Till You Make It: Reputation, Competition, and Yelp Review Fraud. *Management Science 62*(2), 3412–3427.

Malmendier, U. (2016). The Bidder's Curse: Reply. *American Economic Review 106*(4), 1195–1213.

Malmendier, U. and Y. H. Lee (2011). The Bidder's Curse. *American Economic Review 101*(2), 749–787.

Mas, A. (2003). Pay, reference points, and police performance. *Quarterly Journal of Economics 121*(3), 783–821.

Mayzlin, B. D., Y. Dover, and J. Chevalier (2014). Promotional Reviews: An Empirical Investigation of Online Review Manipulation. *American Economic Review 104*(8), 2421–2455.

Melnik, M. I. and J. Alm (2002). Does a Seller's Ecommerce Reputation Matter? Evidence from eBay Auctions. *The Journal of Industrial Economics 50*(3), 337–349.

Mikians, J., L. Gyarmati, V. Erramilli, and N. Laoutaris (2012). Detecting price and search discrimination on the Internet. *Proceedings of the 11th ACM Workshop on Hot Topics in Networks acn*, 79–84.

Nosko, C. and S. Tadelis (2015). The Limits of Reputation in Platform Markets: An Empirical Analysis and Field Experiment. *NBER Working Paper 20830*.

Resnick, P., R. Zeckhauser, J. Swanson, and K. Lockwood (2006). The value of reputation on eBay: A controlled experiment. *Experimental Economics 9*(2), 79–101.

Rotemberg, J. J. (2011). Fair Pricing. *Journal of the European Economic Association 9*(5), 952–981.

Roth, A. E. (2007). Repugnance as a Constraint on Markets. *Journal of Economic Perspectives 21*(3), 37–58.

Schneider, H. S. (2016). The Bidder's Curse: Comment. *American Economic Review 106*(4), 1182–1194.

Shafir, E., P. Diamond, and A. Tversky (1997). Money Illusion. *Quarterly Journal of Economics 112*(2), 341–373.

Shiller, B. R. (2014). First-Degree Price Discrimination Using Big Data. *mimeo*.

Simonsohn, U. and D. Ariely (2008). When Rational Sellers Face Nonrational Buyers: Evidence from Herding on eBay. *Management Science 54*(9), 1624–1637.

Simonsohn, U. and G. Loewenstein (2006). Mistake #37: The effect of previously encountered prices on current housing demand. *Economic Journal 116*(508), 175–199.

Tadelis, S. (2016). Reputation and Feedback Systems in Online Platform Markets. *Annual Review of Economics 8*(1), 321–340.

Turow, J., L. Feldman, and K. Meltzer (2005). Open to Exploitation: America's Shoppers Online and Offline. *Report of the Annenberg Public Policy Center of the University of Pennsylvania*.

Wang, T., E. Gal-Or, and R. Chatterjee (2009). The Name-Your-Own-Price Channel in the Travel Industry: An Analytical Exploration. *Management Science 55*(6), 968–979.

Ward, M. (2000). Amazon's old customers' pay more. *BBC News 8.9.2000*.

Weaver, R. and S. Frederick (2012). A Reference Price Theory of the Endowment Effect. *Journal of Marketing Research 49*(5), 696–707.

Xia, L., K. B. Monroe, and J. L. Cox (2004). The Price Is Unfair! A Conceptual Framework of Price Fairness Perceptions. *Journal of Marketing 68*(4), 1–15.

Yi, S. and H. Baumgartner (2004). Coping With Negative Emotions in Purchase-Related Situations. *Journal of Consumer Psychology 14*(3), 303–317.

Appendix C: Non-linear regression results (not for print)

In the following, I present the marginal effects obtained from the Probit equivalents of the linear probability models (LPM) which have been reported before. For continuous variables, this corresponds to the average effect of the derivative of the likelihood function with respect to that variable. For categorical variables, this is the average of the difference between the predicted probabilities at that variables' respective values. I also estimated all models via Poisson-regressions. The corresponding marginal effect are not reported here but yield the same qualitative, and in most cases also the same quantitative, results. Table 6 below reports the marginal effects of the results when regression model (5) is estimated by Probit. These results therefore correspond to the results of the LPM reported in table 5.

	(1)	(2)	(3)	(4)
SixDayBin#2	-0.158***	-0.152***	-0.143***	-0.136***
	(0.040)	(0.039)	(0.036)	(0.036)
SixDayBin#3	-0.171***	-0.166***	-0.149***	-0.143***
	(0.041)	(0.040)	(0.038)	(0.037)
SixDayBin#4	-0.144***	-0.138***	-0.119***	-0.112***
	(0.044)	(0.044)	(0.041)	(0.040)
SixDayBin#5	-0.202***	-0.197***	-0.179***	-0.173***
	(0.041)	(0.040)	(0.039)	(0.039)
SixDayBin#6	-0.160***	-0.149***	-0.161***	-0.148***
	(0.049)	(0.049)	(0.042)	(0.042)
Buyer's Score >10		-0.023		-0.017
		(0.021)		(0.020)
Buyer's Score >100		-0.035		-0.025
		(0.021)		(0.020)
Buyer's Score >1000		0.061		0.083*
		(0.046)		(0.047)
First reviews only	no	no	yes	yes
Observations	3,575	3,550	2,283	2,265

Table 6. Marginal effects from Probit estimates of regressing an indicator for comment in the review which refer to the seller's sales strategy on dummies for the six-day-bin since the auction when the review was written and the buyers' own feedback score. Columns 1 & 2 use all collected reviews for the auction and report standard errors clustered on the buyer level. Columns 3 & 4 only use the first review a buyer posted for the auction and report robust standard errors. Significance levels: *** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table 7 on the next page reports marginal effects of the results when regression model (3) is estimated by Probit. These results therefore correspond to the results of the LPM reported in table 2. Also on the same page, table 8 reports marginal effects of a Probit model which correspond to the results of the LPM reported in table 3.

	(1)	(2)	(3)	(4)
Auction	0.102***	0.102***	0.100***	0.100***
	(0.008)	(0.008)	(0.007)	(0.007)
Buyer's Score >10		-0.018**		-0.015***
		(0.007)		(0.006)
Buyer's Score >100		-0.015**		-0.012*
		(0.008)		(0.006)
Buyer's Score >1000		-0.004		0.002
		(0.014)		(0.013)
First reviews only	no	no	yes	yes
Observations	18,750	18,594	15,857	15,721

Table 7. Marginal effects from Probit estimates of regressing an indicator for non-positive feedback on a dummy for whether the review was for the auction and dummies for the reviewing buyers' own feedback score. Columns 1 & 2 use all collected reviews and report standard errors clustered on the buyer level. Columns 3 & 4 only use the first review a buyer posted and report robust standard errors. Significance levels: *** $p<0.01$, ** $p<0.05$, * $p<0.1$

	(1)	(2)	(3)	(4)
Δ_{10}Price	0.022***	0.022***	0.023***	0.023***
	(0.005)	(0.005)	(0.005)	(0.005)
Buyer's Score >10		-0.036		-0.034
		(0.029)		(0.025)
Buyer's Score >100		-0.045		-0.040
		(0.029)		(0.025)
Buyer's Score >1000		0.011		0.035
		(0.052)		(0.050)
First reviews only	no	no	yes	yes
Observations	3,575	3,550	2,283	2,265

Table 8. Marginal effects from Probit estimates of regressing an indicator for non-positive feedback on the difference between the auction price and the fixed price divided by 10 and dummies for the reviewing buyers' own feedback score. Columns 1 & 2 use all collected reviews and report standard errors clustered on the buyer level. Columns 3 & 4 only use the first review a buyer posted and report robust standard errors. Significance levels: *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

In table 9 on the next page, I report the marginal effect of estimating the interaction model (9) by Probit. The results of interaction term in non-linear models such as Probit, are sensitive to how they are computed, in particular what the relevant covariates are.[17] The approach use here is the following: Model (4) is estimated as a Probit-model. The average marginal effects of the coefficient on $\Delta_{10}Price_i$ is then computed for each of the six possible configuration when either one or none of the $SixDayBin\#t_i$-dummies is set to, i.e. for all six-day-bins. The so-determined six price slopes therefore correspond to the estimated coefficient $\hat{\beta}_1$ in the corresponding table 4 for the first six-day-bin and to the sum of the corresponding coefficients $\hat{\beta}_1 + \hat{\beta}_t$ from that table for the later six-day-bins. Therefore, the significance test on these marginal effects have the same underlying null hypothesis as the F-test reported in table 4.

[17] See Ai, C., and Norton, E. C. (2003). Interaction terms in logit and probit models. *Economics Letters 80*, 123-129. For an accessible guide on how to compute marginal effects for interaction terms see Karaca-Mandic, P., Norton, E. C. and Dowd, B. (2011). Interaction Terms in Nonlinear Models. *Health Services Research, 47*(1), 255-274.

	(1)	(2)	(3)	(4)
Δ_{10}Price (at SixDayBin#1)	0.125***	0.125***	0.116***	0.116***
	(0.023)	(0.024)	(0.027)	(0.027)
Δ_{10}Price (at SixDayBin#2)	0.008	0.009	0.008	0.008
	(0.006)	(0.006)	(0.006)	(0.006)
Δ_{10}Price (at SixDayBin#3)	0.009	0.007	0.006	0.007
	(0.009)	(0.009)	(0.010)	(0.009)
Δ_{10}Price (at SixDayBin#4)	0.020	0.021	0.020	0.020
	(0.019)	(0.019)	(0.018)	(0.018)
Δ_{10}Price (at SixDayBin#5)	-0.004	-0.003	-0.036	-0.035
	(0.016)	(0.016)	(0.023)	(0.022)
Δ_{10}Price (at SixDayBin#6)	0.033	0.028	0.040*	0.035
	(0.025)	(0.026)	(0.023)	(0.024)
SixDayBin#2	-0.107**	-0.105*	-0.131**	-0.126**
	(0.054)	(0.054)	(0.058)	(0.057)
SixDayBin#3	-0.118**	-0.116**	-0.129**	-0.128**
	(0.056)	(0.043)	(0.060)	(0.059)
SixDayBin#4	-0.078	-0.075	-0.076	-0.072
	(0.062)	(0.061)	(0.066)	(0.065)
SixDayBin#5	-0.129**	-0.128**	-0.092	-0.090
	(0.062)	(0.061)	(0.072)	(0.071)
SixDayBin#6	-0.075	-0.069	-0.136**	-0.125*
	(0.073)	(0.073)	(0.066)	(0.066)
Buyer's Score >10		-0.031		-0.031
		(0.028)		(0.024)
Buyer's Score >100		-0.047*		-0.039
		(0.028)		(0.024)
Buyer's Score >1000		0.019		0.028
		(0.051)		(0.050)
First reviews only	no	no	yes	yes
Observations	3,575	3,550	2,283	2,265

Table 9. Marginal effects from Probit estimates of regressing estimates of regressing an indicator for non-positive feedback on the difference between the auction price and the fixed price divided by 10, a dummy for the six-day-bin since the auction in which the review was written, its interaction with the price variable, and dummies for the reviewing buyers' own feedback score. Reported marginal effects refer to the estimated price slope and the corresponding shift of the intercept when $SixDayBin\#t = 1$ with $t \in \{1, .., 6\}$ and all other six-day-dummies are zero. Columns 1 & 2 use all collected reviews and report standard errors clustered on the buyer level. Columns 3 & 4 only use the first review a buyer posted and report robust standard errors. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.